Wax

ARPA Network Information Center
Stanford Research Institute
Menlo Park, California 94025
NIC #30797
Packet Radio Temporary Note #109

Network Analysis Corp.
May, 1974

A SIMULATION OF THE PACKET RADIO NETWORK

1. INTRODUCTION

A simulation program of a packet radio network is described. The main features which distinguish this network from a point-to-point (PTP) packet switching network (e.g., the ARPANET) are: (i) devices in the network transmit packets by using a random access scheme, and (ii) devices broadcast so that signals can be received by several devices simultaneously. We begin by, briefly, indicating the implication of the above characteristics, state the objectives of the simulation, and then proceed with the description of the program.

When using a random access scheme, there is a possibility that more than one packet are simultaneously received by a receiver due to independent transmissions of several devices; in that event, none of these packets are correctly received, and the corresponding devices have to retransmit their packets. This implies that, unlike in a PTP network, the probability of error due to overlapping packets is much higher than the probability of error due to other causes such as Gausian or impulse noise. Furthermore, the probability of error varies widely depending on the amount of traffic on the channel. The broadcast nature of the network implies that there is correlation between the probabilities of successful transmission in different parts of the network.

The communication system simulated contains three types of devices, terminals, repeaters, and stations. The station is considered as an interface communication device of the broadcast network to a higher level network or to a computer installation. Terminals are considered as traffic sources and traffic sinks, they transmit packets to the station and receive response packets from the station; they may be mobile. The basic function of the repeater

is to extend the effective range of the terminals and the station. The network simulated contains repeaters and stations which are placed at fixed locations chosen at random, as shown in Figure 1. Terminals originate at random times and are placed in random locations on the plane. A terminal is considered to depart from the system once it completes its communication. A more detailed description of the system simulated can be found in [NAC (a)], and [NAC (b)].

The program was coded in FORTRAN. It includes a total of 33 subroutines, approximately 3,000 statements. The compilation time on a CDC 6600 takes approximately 20 CP sec, and the running time for meaningful results takes from 100 - 300 CP sec. The storage requirement of the program is $245{,}000_8$ 60 bit words.

The simulator has already been used for improving communication protocols, and to answer questions related to the trade offs between device range and device interference, and the trade offs between a single and a dual data rate system [NAC (c)].

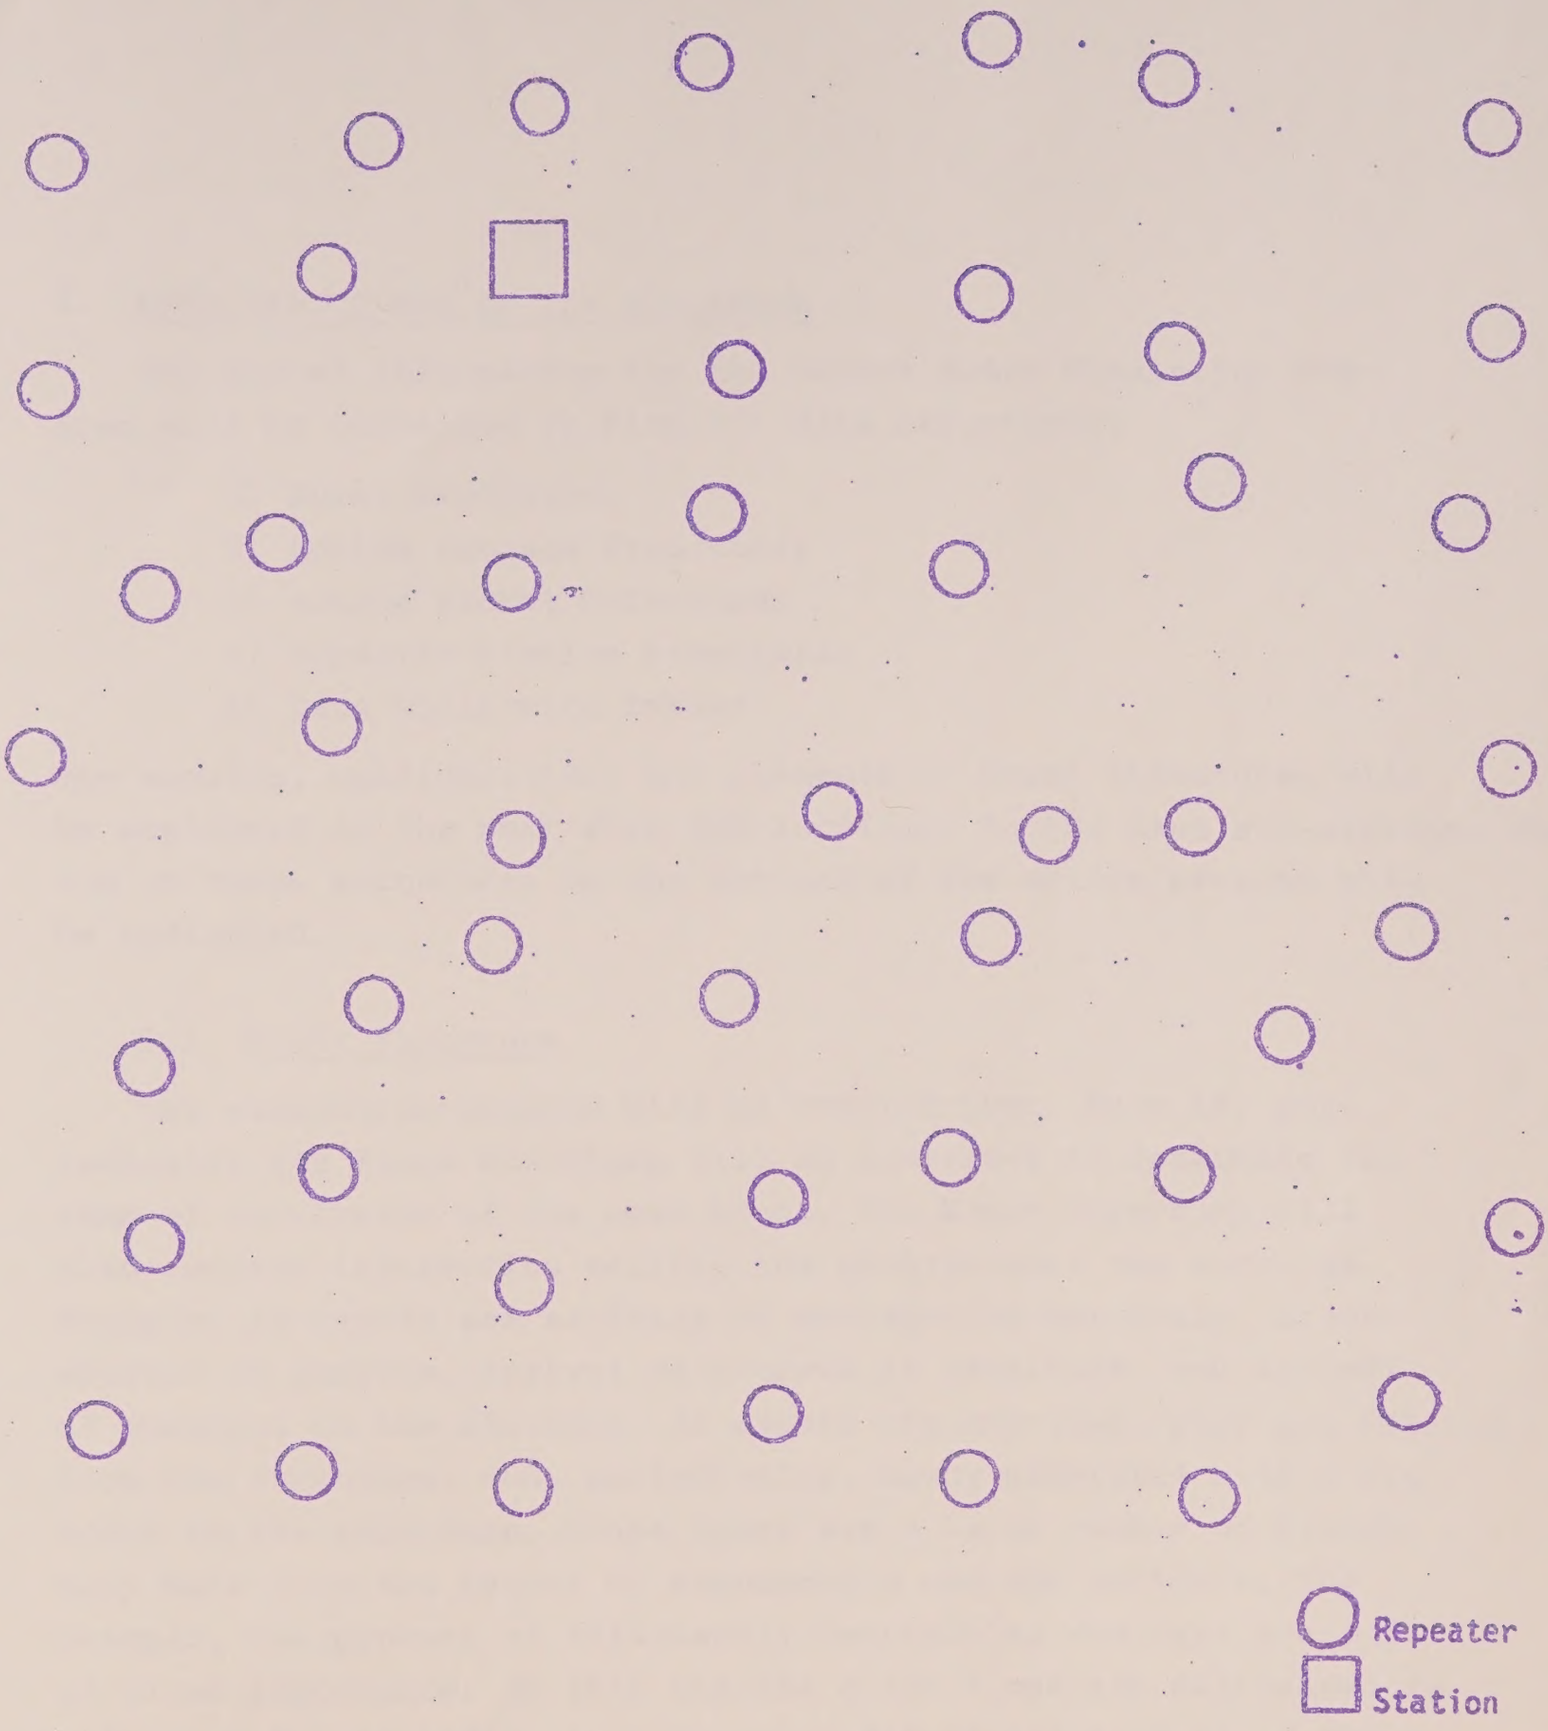

Figure 1 Device Locations.

2. DATA STRUCTURES OF THE SIMULATOR

The global information for the Packet Radio Simulation Program will be contained in five (5) data structures:

1) Event Structure;
2) Active Message Structure;
3) Active Packet Structure;
4) Repeater-Station Structure;
5) Data Collection Tables.

The meaning, configuration, and elements of these structures will be explained in the next five sub-sections. In the last sub-section, the use of these structures in the context of the entire program will be indicated.

2.1 Event Structure

The simulation program will be event-driven. That is, periodically the Event Structure will be consulted to determine the time of occurrence of the next event. The Event Structure will also contain information telling the program what the event is. Examples of events are arrivals of messages to terminals, transmission of packets, arrival of packets at receivers, and arrival of messages to the stations. As events are executed, they are deleted from the structure; and, periodically, newly-generated events are added to the structure. Since there are a large number of events; many more than the number of exongoneous message arrivals, for example, the process of efficiently determining the next event is of vital importance. To this end the event times are maintained in a "heap." Corresponding to each event "i" is its time "t_i," the device subroutine "d_i" to which it refers (e.g. station, terminal, repeater), the index "w_i" of the device in question (i.e. which repeater, which station, etc.), a number representing the

point at which the routine is entered, and, finally, the packet number of the packet in question. In addition, there is the heap index vector itself. "h_j" points to the event which occupies the jth position in the heap. A heap is a structure in which $t_{hj} \leq \text{Max } \{t_{h2j}, t_{h_{2j+1}}\}$. At all times t_{h_1} is the smallest t_i over all the events. This structure allows for quite rapid selection of the minimum t_i while using a minimal amount of storage. [Knuth; 1973], [Kershenbaum and Van Slyke; 1972, Appendix], [Williams; 1964]. In order to efficiently eliminate old events and to reuse the space created by the elimination, a garbage stack of unused locations is maintained.

2.2 Active Message Structure

The external or exogoneous traffic which flows into the Packet Radio Network consists of messages which arrive at terminals or stations and represent information that users wish to send to a location using the net. Messages are to be distinguished from packets which carry the message information internally in the net. In general, there may be several packets in the network carrying copies of the same message or of parts of that message. Messages are added to the Active Message Structure when the event corresponding to their generation occurs. It stays on the list until the last packet containing the message is dealt with. Associated with each message "i" is its arrival time "t_i;" the current number of packets representing the message "n_i" (When "n_i" is reduced to zero (0), the message is removed from the structure.); its length "l_i;" "x_i" and "y_i", the co-ordinates of the terminal or station originating the message; and other pieces of data, such as the repeater with which the terminal communicates, the total number of packets, and arrival time of message, which may be needed for output statistics. In order to eliminate old messages and add new ones efficiently, the messages are kept in a doubly-linked list structure. There is a list kept for messages and one for unused message spaces. This information is kept in two vectors, "f" and "b". Thus, "f_i" repre-

Digitized by the Internet Archive
in 2025 with funding from
Amateur Radio Digital Communications, Grant 151

https://archive.org/details/temporary-note-109

sents whatever follows the space in question ("i"), be it message or empty space; while "b_i" represents whatever precedes the space. The unit being followed or preceded may be an empty space or may contain a message.

2.3 Active Packet Structure

The active packets are kept in a list. Associated with the ith packet is a pointer to its corresponding message and several words which state the variable part of the packet label for use by the routing algorithms. A garbage stack similar to the one used for events keeps track of vacant spaces in the Packet Structure.

2.4 Repeater-Station Structure

The Repeater-Station Structure is basically a list of the repeaters and stations, their locations, their possible neighbors (repeaters and stations within range of their transmitters and receivers), and their state. Associated with the ith station or repeater are "x_i" and "y_i", the co-ordinates of the device; and a state vector which, among others, specifies items such as whether the repeater or station is busy, free, turned off, or failed; and whether the device is a station or a repeater. In addition, there is a list of repeaters and stations adjacent to the ith repeater or station.

2.5 Data Collection Tables

Sufficient statistics for the evaluation of the system performance measures (see Section 3.4) are kept in the data collection tables.

2.6 An Outline of the Use of Data Structures

The global simulation structure is based on events of packets arriving at devices: stations, terminals, and repeaters. Figure 2.

schematically illustrates the program flow. It is quite simplified, one simplification being especially important and needing emphasis. The events of retransmissions and acknowledgements resulting from a packet's arriving at a repeater, station, or terminal are not necessarily generated immediately, but may depend on the arrival of packets at the device subsequent to the packet which gives rise to the new events.

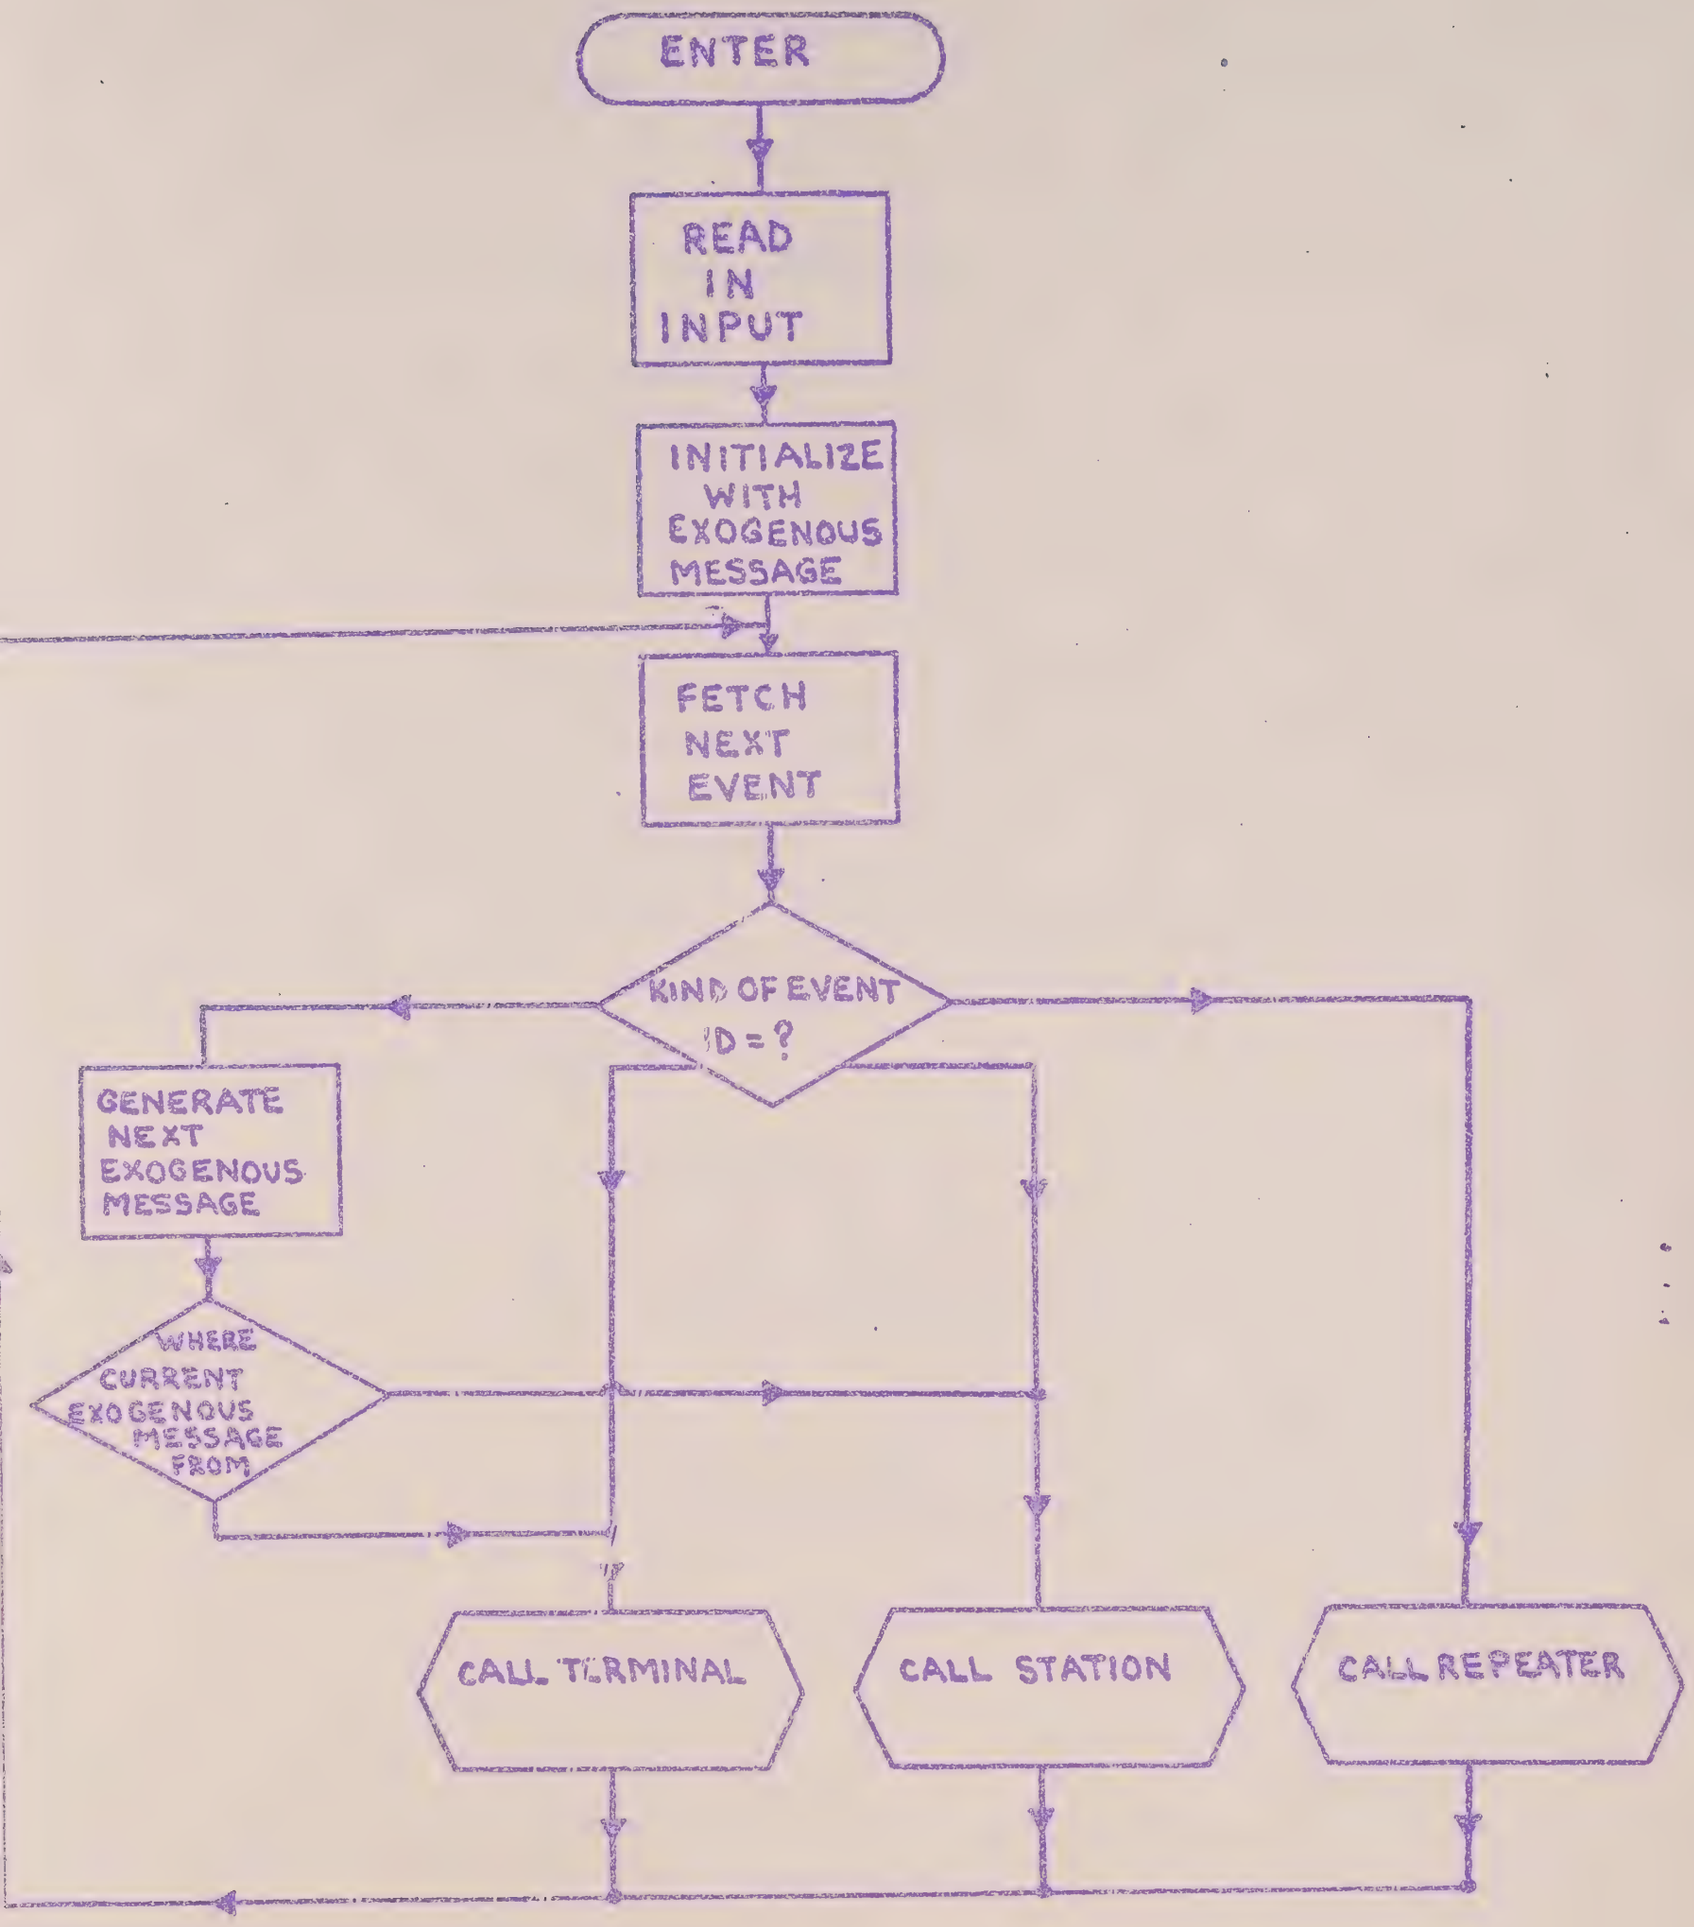
ENTER
READ IN INPUT
INITIALIZE WITH EXOGENOUS MESSAGE
FETCH NEXT EVENT
KIND OF EVENT ID = ?
GENERATE NEXT EXOGENOUS MESSAGE
WHERE CURRENT EXOGENOUS MESSAGE FROM
CALL TERMINAL
CALL STATION
CALL REPEATER

Figure 2

3. THE COMMUNICATION SYSTEM SIMULATED

3.1 General System Description

Channel and Access Mode: The communication channel is shared for transmission in both directions, to the station or from the station to terminals. The channel access mode is the Non-Persistent Carrier Sense Multiple Access (CSMA) [Kleinrock, 1974]. That is, when a packet is ready for transmission, the device senses the channel and transmits the entire packet if the channel is idle. If the channel is busy, the device reschedules the packet for some future random time at which it senses the channel again, and the procedure repeated.

Capture: A zero capture system is simulated. That is, whenever the reception of more than one packet overlap in time, none of the packets is correctly received.

Packet Types Simulated:

IP - Information Packet

ETE - ETE Ack Short Packet (Assumed to be 10% of the length of an information packet).

SP - Search Packet, transmitted by Terminal or Repeater to all devices and aimed to identify a specific receiver (short packet).

RSP - Response to a Search Packet, transmitted by Repeater or Station which received a SP and is available for handling packets. This packet contains the label of the transmitting device and is addressed to all devices.

3.2 Routing Algorithms

Three routing algorithms which are implemented in the simulator are briefly described in the following paragraphs. A detailed description of the routing schemes is given in [NAC (b)].

Hierarchical Labeling: The hierarchical labeling routing scheme enables point-to-point routing between devices along an "efficient path". It is obtained by assigning to every repeater a label, which forms, functionally, a hierarchical structure (see Figure 3). The label assigned contains the following information:

a. A specific address of the repeater for routing purposes.

b. The minimum number of hops to the nearest station.

c. The specific address of all repeaters on the shortest path to the station, and the address of the repeater to which a packet has to be transmitted when destined to the station.

In the hierarchical labeling algorithm an information packet (IP) is addressed to one device. If it is received by a device to which it is not addressed, then the receiving device is closer to the destination than the device to which the packet was addressed. If the preassigned path is temporarily blocked, the packet may depart from it. It then uses the most efficient path from its new location.

Directed Routing (one Level Labels): Directed routing is a simplified version of the hierarchical labeling algorithm in which the only information preserved is the direction TO or FROM station. Repeaters are assigned labels that indicate the hierarchy level, or the number of hops to the nearest station. When a device transmits

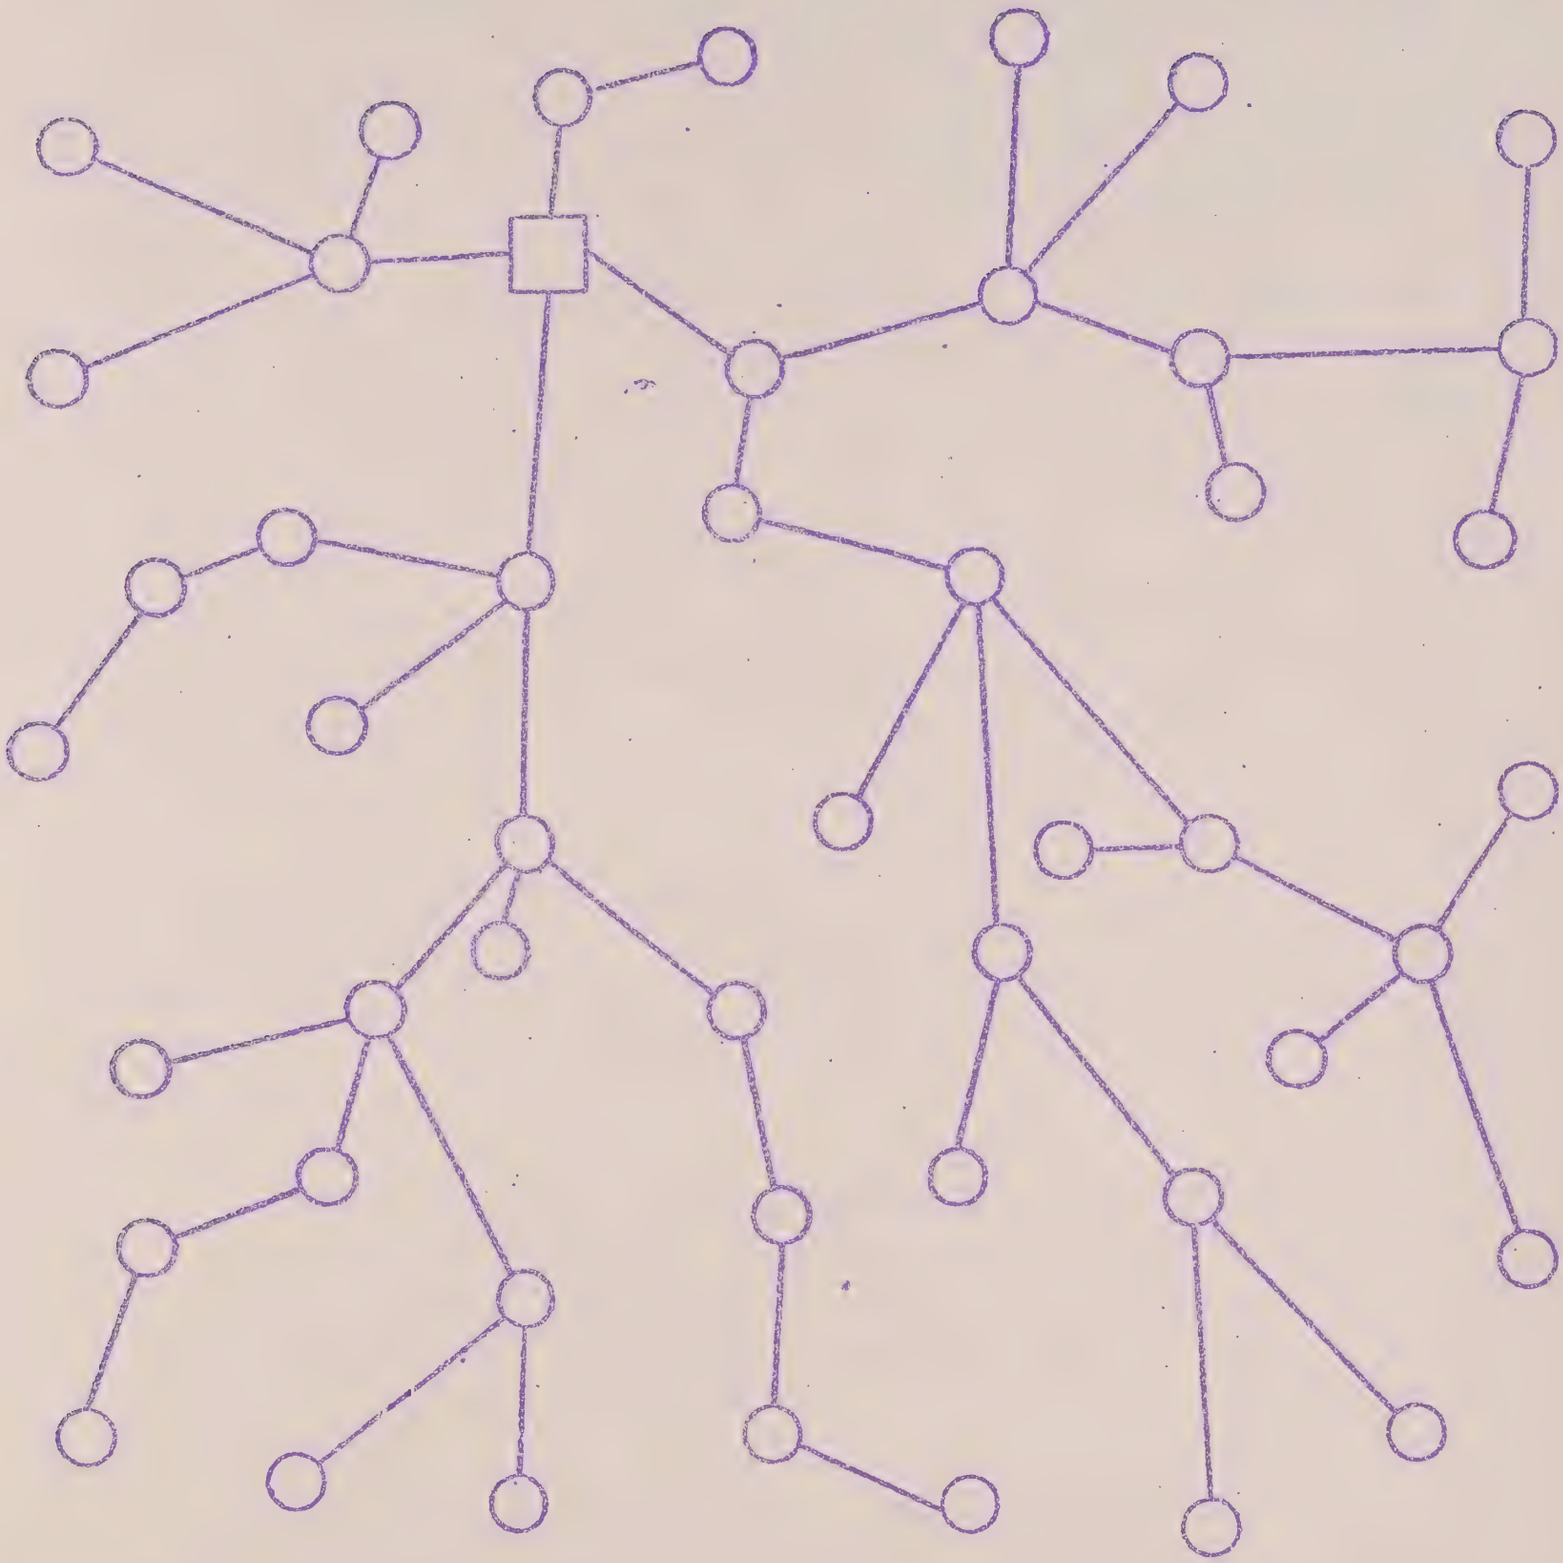

Figure 3

a packet, the packet is addressed to all repeaters (stations) that are closer to the destination than the transmitting device. Many devices can receive the same packet and it may arrive at more than one station. The acknowledgement schemes, the station-station protocol, and the station-terminal protocol must then resolve this problem.

Undirected Routing Plus Repeater Memory: In this algorithm, there is no directionality of transmission. A packet is addressed to all devices that can hear it. To control the problems of cycling and looping, repeaters are assigned storage for unique identifiers of packets that they recently repeated. When a packet is received by a repeater, it compares its identifier with those stored and discards the packet if a match occurs. A maximum handover number (MHN) in the packet will prevent it from being propagated for very long distances. This feature is also used in the other routing algorithms.

3.3 Acknowledgement Schemes

The acknowledgement scheme has particular significance in this system because of the broadcast feature and the limited capability of the repeater for processing and storage. The following acknowledgements are used:

a. End-to-end acknowledgement (ETE Ack) between station and terminal to ensure message integrity. The frequency and precise meaning of the Ack depends on the particular protocol used, and is part of the protocol.

b. A hop-by-hop passive echo acknowledgement (HBH Echo Ack) along the path. When device I transmits a packet, it waits a sufficient time to

allow devices that receive the packet to repeat it. When any of these repeats the packet and the packet is received by device I, it considers it as an Ack. The Ack test includes a unique identifier of the packet and the maximum handover number (MHN). If the receiving device is an end device (terminal or station) it sets the MHN to zero and repeats the header to Ack devices that store the packet. Note that the HBH Echo Ack is nondirected and independent of the path travelled. This Ack simplifies the repeater program since it need not contruct a specific acknowledgement for each packet it correctly receives.

3.4 Performance Measures

A. Throughput

Considering the set of stations and the set of terminals as the end devices, the system throughput (in packets) is defined as the rate of information packets (IP's) that originated at stations and arrived at terminals, plus the rate of IP's that originated at terminals and arrived at stations. Some of the packets contributing to the value of throughput as defined above will be contributed by IP's that have been reactivated. Thus, one can speak of gross throughput as defined above and net throughput in which the rate of reactivated packets is not considered. The net throughput can also be defined as the summation of the rates of IP's that have been end-to-end acknowledged. To obtain a detailed observation of system performance, it is necessary to measure the gross and net throughputs. The difference between the two reflects the efficiency (or deficiency) of the routing and protocols.

B. Delays

The following delays are measured:

1. Terminal delay to identify specific repeater.

2. Terminal delay to establish communication with station and to negotiate protocols.

3. End-to-end delay for an IP

4. Terminal interaction delay as a function of the number of IP's transmitted and received. The interaction delay is defined as the time elapsed from terminal origination to departure.

C. Blocking and Loss

When a terminal does not successfully identify a repeater (or station) after transmitting an SP for the maximum number of times specified, it is considered blocked. In addition, under certain conditions, terminals will depart from the system without completing communication. This will contribute to additional system loss. The blocking and loss are measured separately since the former indicates the difficulty in entering the system, whereas the latter reflects on the inefficiency of the routing.

D. Device Performance

1. Probability that the station is busy.

The station is sampled during the simulation and is assumed busy if it is actively receiving

or transmitting; otherwise, it is assumed idle. This measure is an indication of the channel traffic at the station.

2. Successful completions by repeaters. The number of packets that each repeater has successfully switched are counted. This indicates the distribution of load in the network and reflects on the power duty cycle of repeaters.

E. Other Measures

1. The number of terminals in the system, the total number of packets stored in the system, the number of events to be processed; all as a function of time.

2. The complete output includes a detailed description of the flow of significant communication events. In particular, one can trace a specific packet from its origin to its destination, or trace the complete communication of a particular terminal as a function of time, from its origination until departure. This enables a detailed observation of system performance.

4. LOGICAL OPERATION OF DEVICES

4.1 States of Devices

Each device is characterized by a state vector. Some of the state variables will be needed in the physical devices, for example the label, a parameter indicating the maximum number of transmissions, the maximum handover number to be assigned by repeater and station, the state of occupancy of its storage, etc. Other variables are particular to the simulation. The following are examples.

- Operational State of Device

  PR - Passive Receive State: The device is in receive state and does not sense any carrier.

  AR - Active Receive State.

  AT - Active Transmit State.

  ART- Active Receive and Transmit.

When a device is in state AR or ART, it can be receiving several overlapping packets simultaneously. In the program we use a common channel configuration (half duplex). Thus, since carrier sense is used for channel access, the device can change to AT only from PR and to ART only from AT.

- Number of Overlapping Packets

This number is incremented by one whenever the device is in state AR or ART and a new packet begins to arrive; and

decremented when a packet transmission ends. The number of overlapping packets indicates the number of times an end of packet transmission has to occur before the device changes its state to PR.

- End of Busy Period

This time is recorded for the purpose of saving CPU time. The transmission time of a packet is set to the End of Busy Period plus a random time; otherwise the device may be called to transmit a packet several times during its busy period.

4.2 Terminal

When a terminal originates a message, it begins to transmit and retransmit a SP to identify a specific receiver. If it does not identify a specific receiver after a specified number of transmissions, it departs from the system. We say that such a terminal is blocked. When a terminal identifies a specific receiver, it substitutes the label and MHN sent by the receiver into its IP and and begins to transmit its IP. The IP is retransmitted after short waiting periods of time until a HBH Echo Ack is received. At that time the terminal stores the IP for a longer period of waiting after which the IP is reactivated if an ETE Ack is not received.* The terminal is expecting several IP's from the station, each of which is ETE acknowledged by the terminal. When all IP's from the station are received and ETE acknowledged, the terminal departs from the system.

---

* We use the term retransmission when a device waits a relatively short period of time (less than 2 IP slots) and is awaiting a HBH acknowledgement. We say that a packet is reactivated when an end device stores the packet, awaiting an ETE Ack. When a packet is reactivated it goes through the whole process of retransmissions.

4.3 Repeater

A repeater does not distinguish between IPs or ETE Acks, except for their transmission time. When an IP (or ETE Ack) is received by a repeater (addressed to it) and the repeater has available storage, it stores the packet, decrements the MHN, modifies the packet label according to the routing and begins to transmit and retransmit the packet, awaiting an Echo Ack. When an Echo Ack is received, the repeater discards the packet. When a repeater is not successful in transmitting along the "shortest" path it begins to search for an alternate receiver by transmitting SPs. When one is found it transmits the entire packet to it; otherwise, it discards the packet. When a repeater receives an SP it checks whether it has available storage; if it does, it makes one attempt to transmit a RSP and then discards it. When a repeater receives a RSP it tests whether it needs one; if it does, it uses the label, otherwise it discards it. The repeater currently simulated has buffer storage for two packets; one exclussively used for packets directed towards the station and the second for packets toward the terminal. In addition, the repeater can inspect all packets that it receives, which are stored in common arrays in the simulation program. Thus, from a practical viewpoint, buffer storage for three packets per repeater are provided in the simulation program.

4.4 Station

The storage organization in the simulated station is shown in Figure 4.* There are two queues for active packets. Packets in these queues are active in the sense that they are retransmitted

---

* Figure 4 shows only the storage of IP's and ETE Acks for transmission to terminals in the packet radio network.

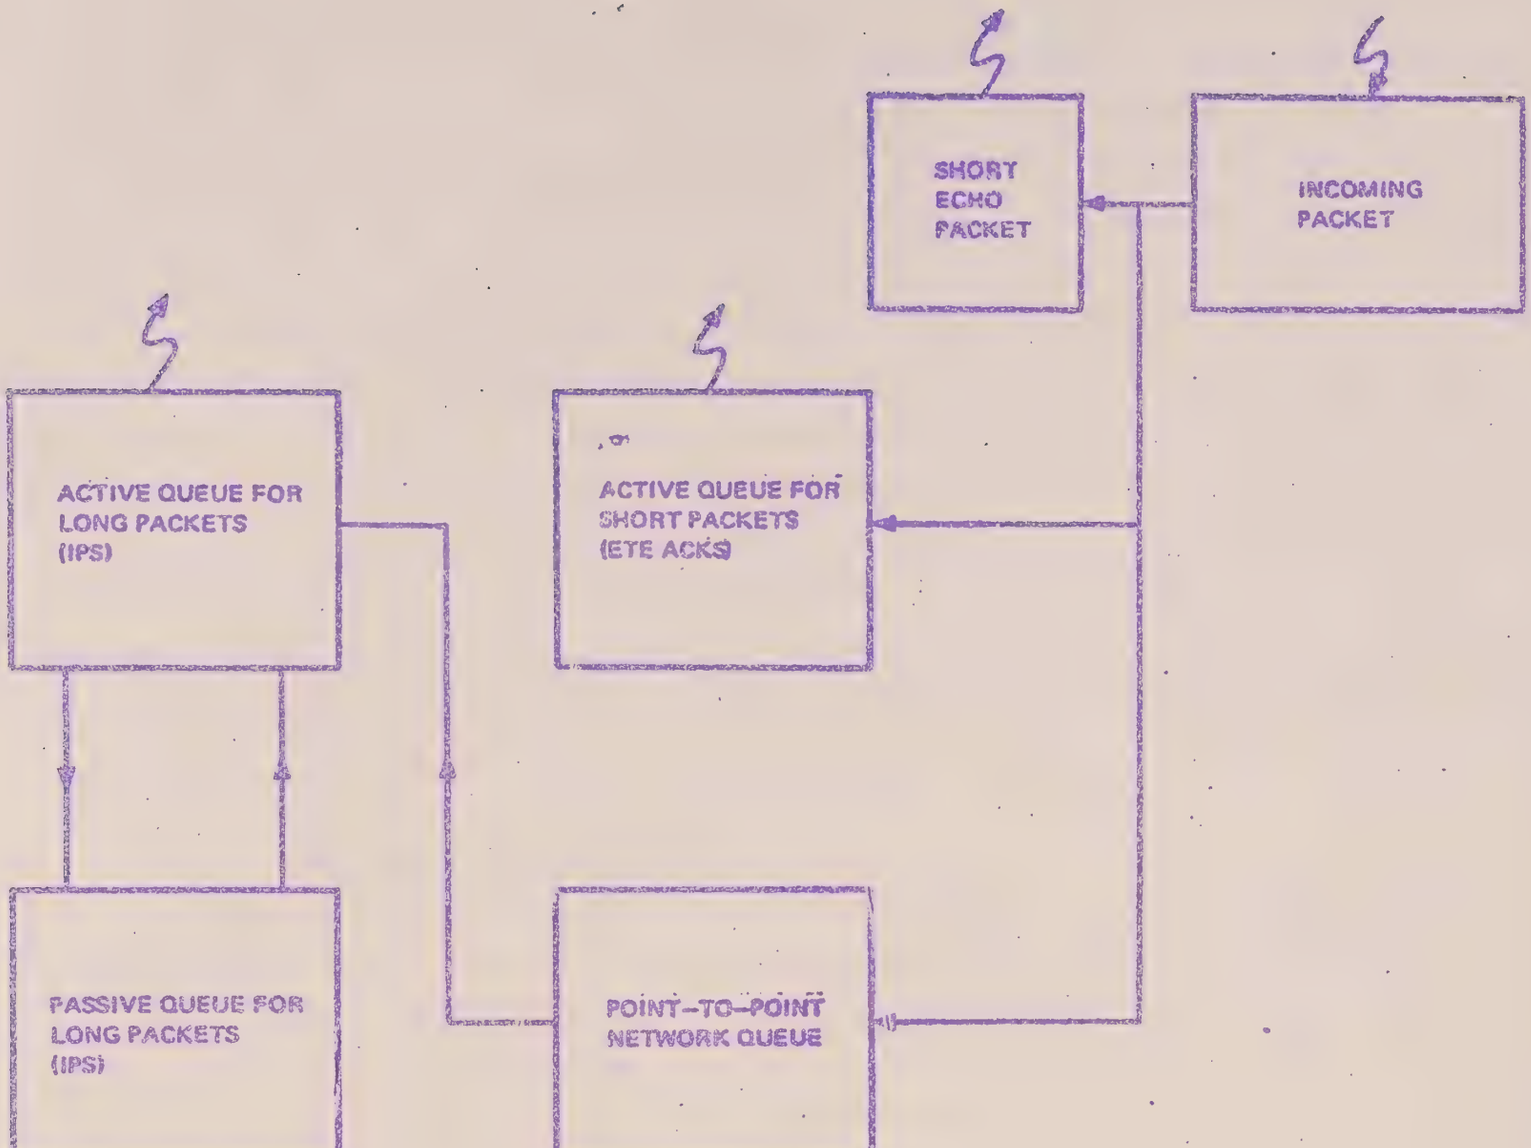

Figure 4 Storage in Station.

after short random periods of waiting until an Echo Ack is received. The active queue for long packets contains IPs from the station to terminals. Once an IP is Echo acknowledged it is stored in the passive queue for a longer period after which it is reactivated if an ETE Ack from the terminal is not received. The active queue for short packets contains ETE Acknowledgement packets to terminals, and these have priority over the long active packets. The ETE Ack packets are, obviously, discarded once an Echo Ack is received. The point-to-point (PTP) network queue simulates the interaction of the packet radio network with a PTP network. When a new IP is received from a terminal, it is stored in the PTP network queue for a random time, after which a response message containing several response IPs to that terminal are generated and placed into the active queue for long packets. The station responds immediately to SPs, and ignores RSPs.

4.5 Flow Diagrams of Devices

Figures 5, 6, and 7 show the flow diagrams of the devices used in the simulator. These diagrams are simplified to the extend that they show "what to do" but not sufficiently detailed to show "how to do it". The latter depends on the particular system simulated, i.e. the routing, the channel configuration, etc.

A device is called from the subroutine EVENT; the calling sequence includes, among others, an interrupt number which indicates to the device the task that it has to perform. The only event which is external to a device is that with an interrupt=1. All other enteries to a device are due to events generated by the device itself. Thus, the number of exogeneous events is very small compared to the number of events generated by devices; in particular when the offered data rate to the system is high.

It is clear that when the offered traffic rate to the system is high, there will be many collisions of packets. Thus to save

computing time, each device was coded so that it does not examine the content of the packet or whether the packet is addressed to it, at the begining of packet reception (see diagrams). This is done at the end of packet reception, providing there was no interference.

There are parts of the subroutines of Repeater, Station, and Terminal which are identical and thus coded into subroutines. These relate to the identification of the header content, and the channel access mode. Some of these can be associated with the modem of the physical devices. The parts which differ involve mainly the processing needed once the packet has been identified, e.g., queueing at the station, ETE Ack by end devices, etc.

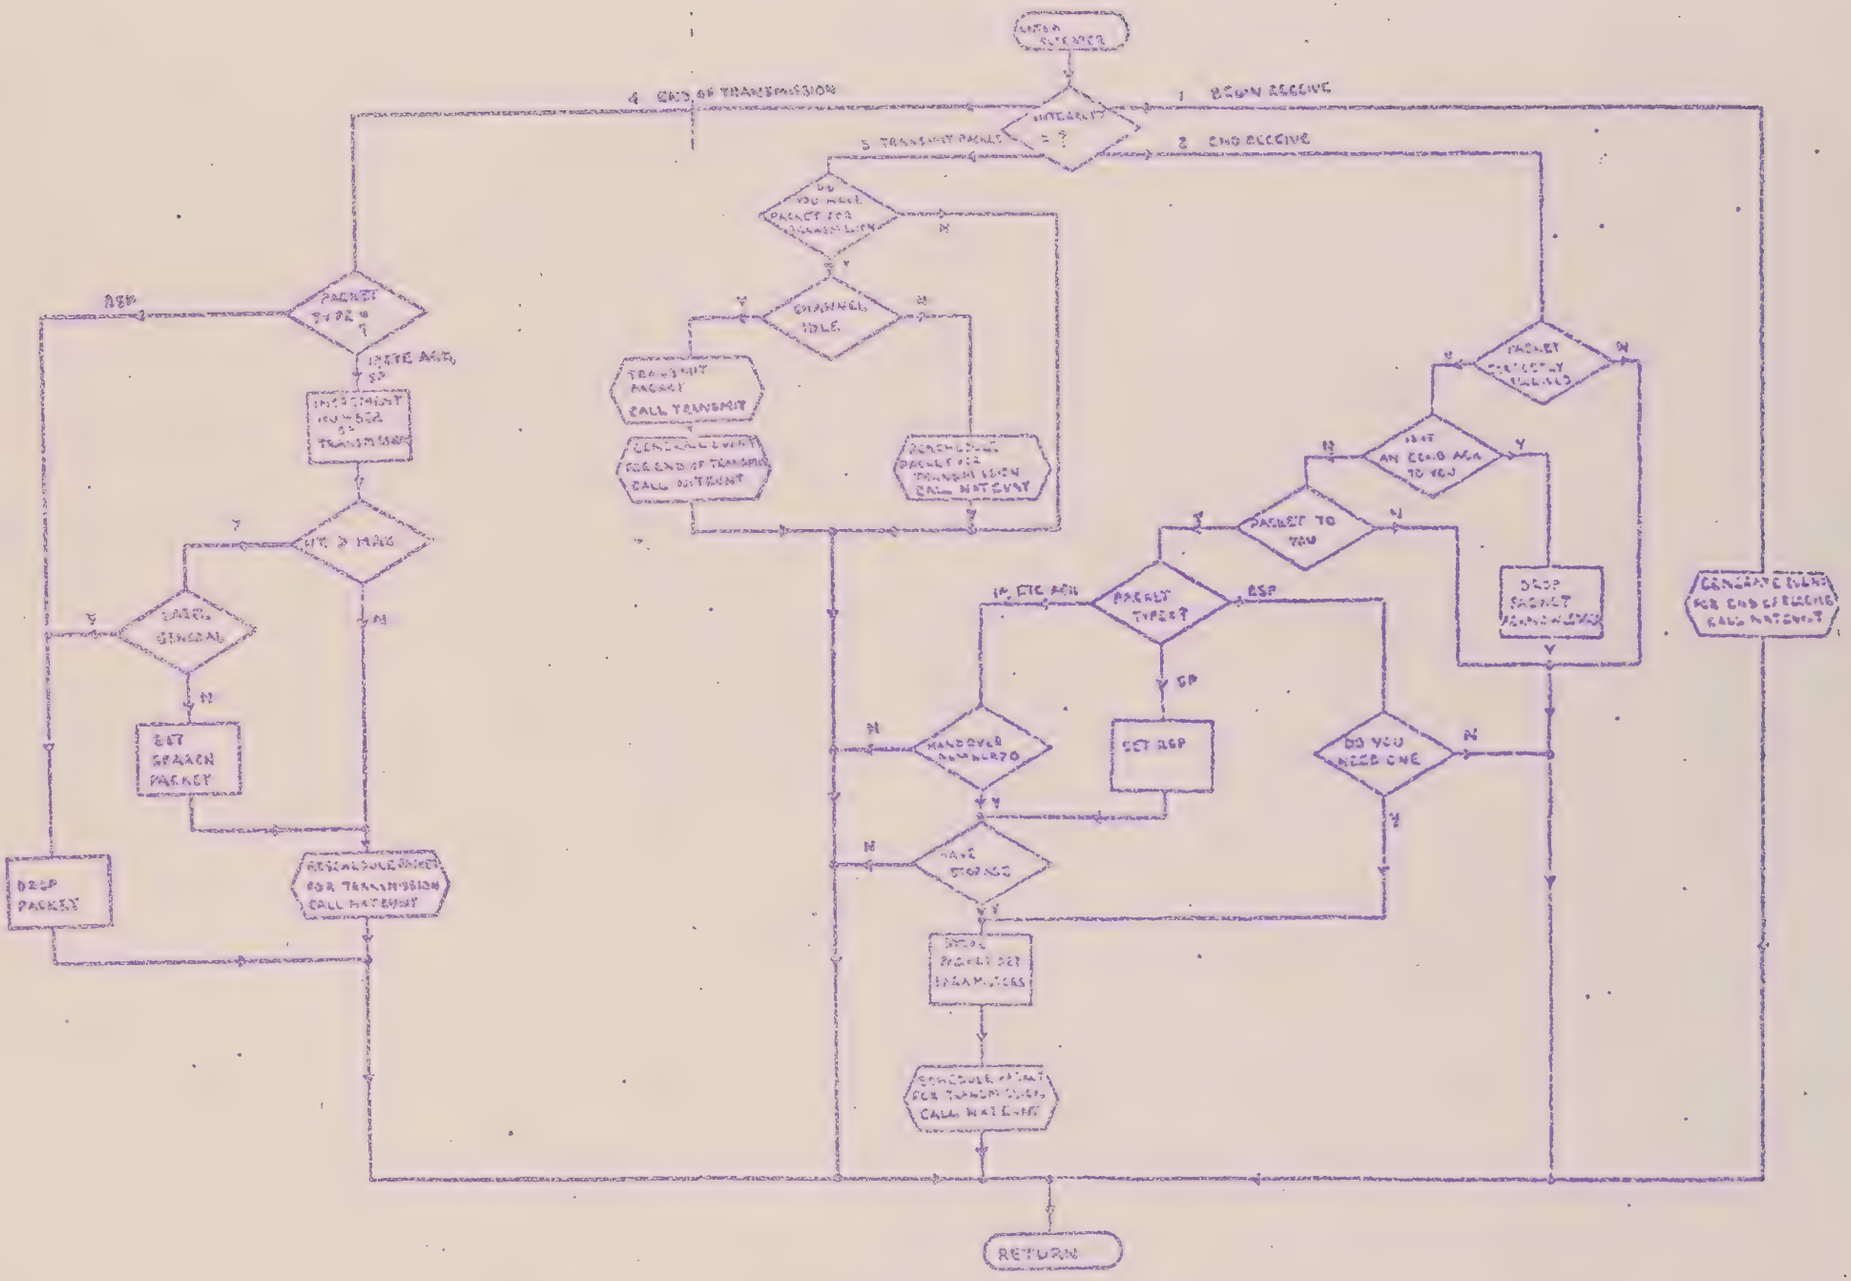
4 END OF TRANSMISSION
1 BEGIN RECEIVE
INTERRUPT = ?
5 TRANSMIT PACKET
2 END RECEIVE
DO YOU NEED PACKET FOR TRANSMISSION
CHANNEL IDLE
TRANSMIT PACKET
CALL TRANSMIT
CANCEL CALL EVENT FOR END OF TRANSMIT CALL INTERRUPT
RESCHEDULE PACKET FOR TRANSMISSION CALL NATEVNT
PACKET TYPE = ?
RSP
INTR
INITIATE ACK, SP
INITIATE NEXT PACKET FOR TRANSMISSION
NT > MAX
CATCH = GENERAL
SET SEARCH PACKET
DROP PACKET
RESCHEDULE PACKET FOR TRANSMISSION CALL NATEVNT
PACKET CORRECTLY RECEIVED
N
IS IT AN ECHO ACK TO YOU
Y
PACKET TO YOU
N
DROP PACKET
IA ETC ACK
PACKET TYPE?
RSP
SP
HANDOVER BOUNDARY?
SET RSP
DO YOU NEED ONE
N
Y
HAVE STORED
INITIATE PACKET SET TRANSMISSION
SCHEDULE INTR FOR TRANSMISSION CALL NATEVNT
CANCEL EVENT FOR END OF RECEIVE CALL NATEVNT
RETURN

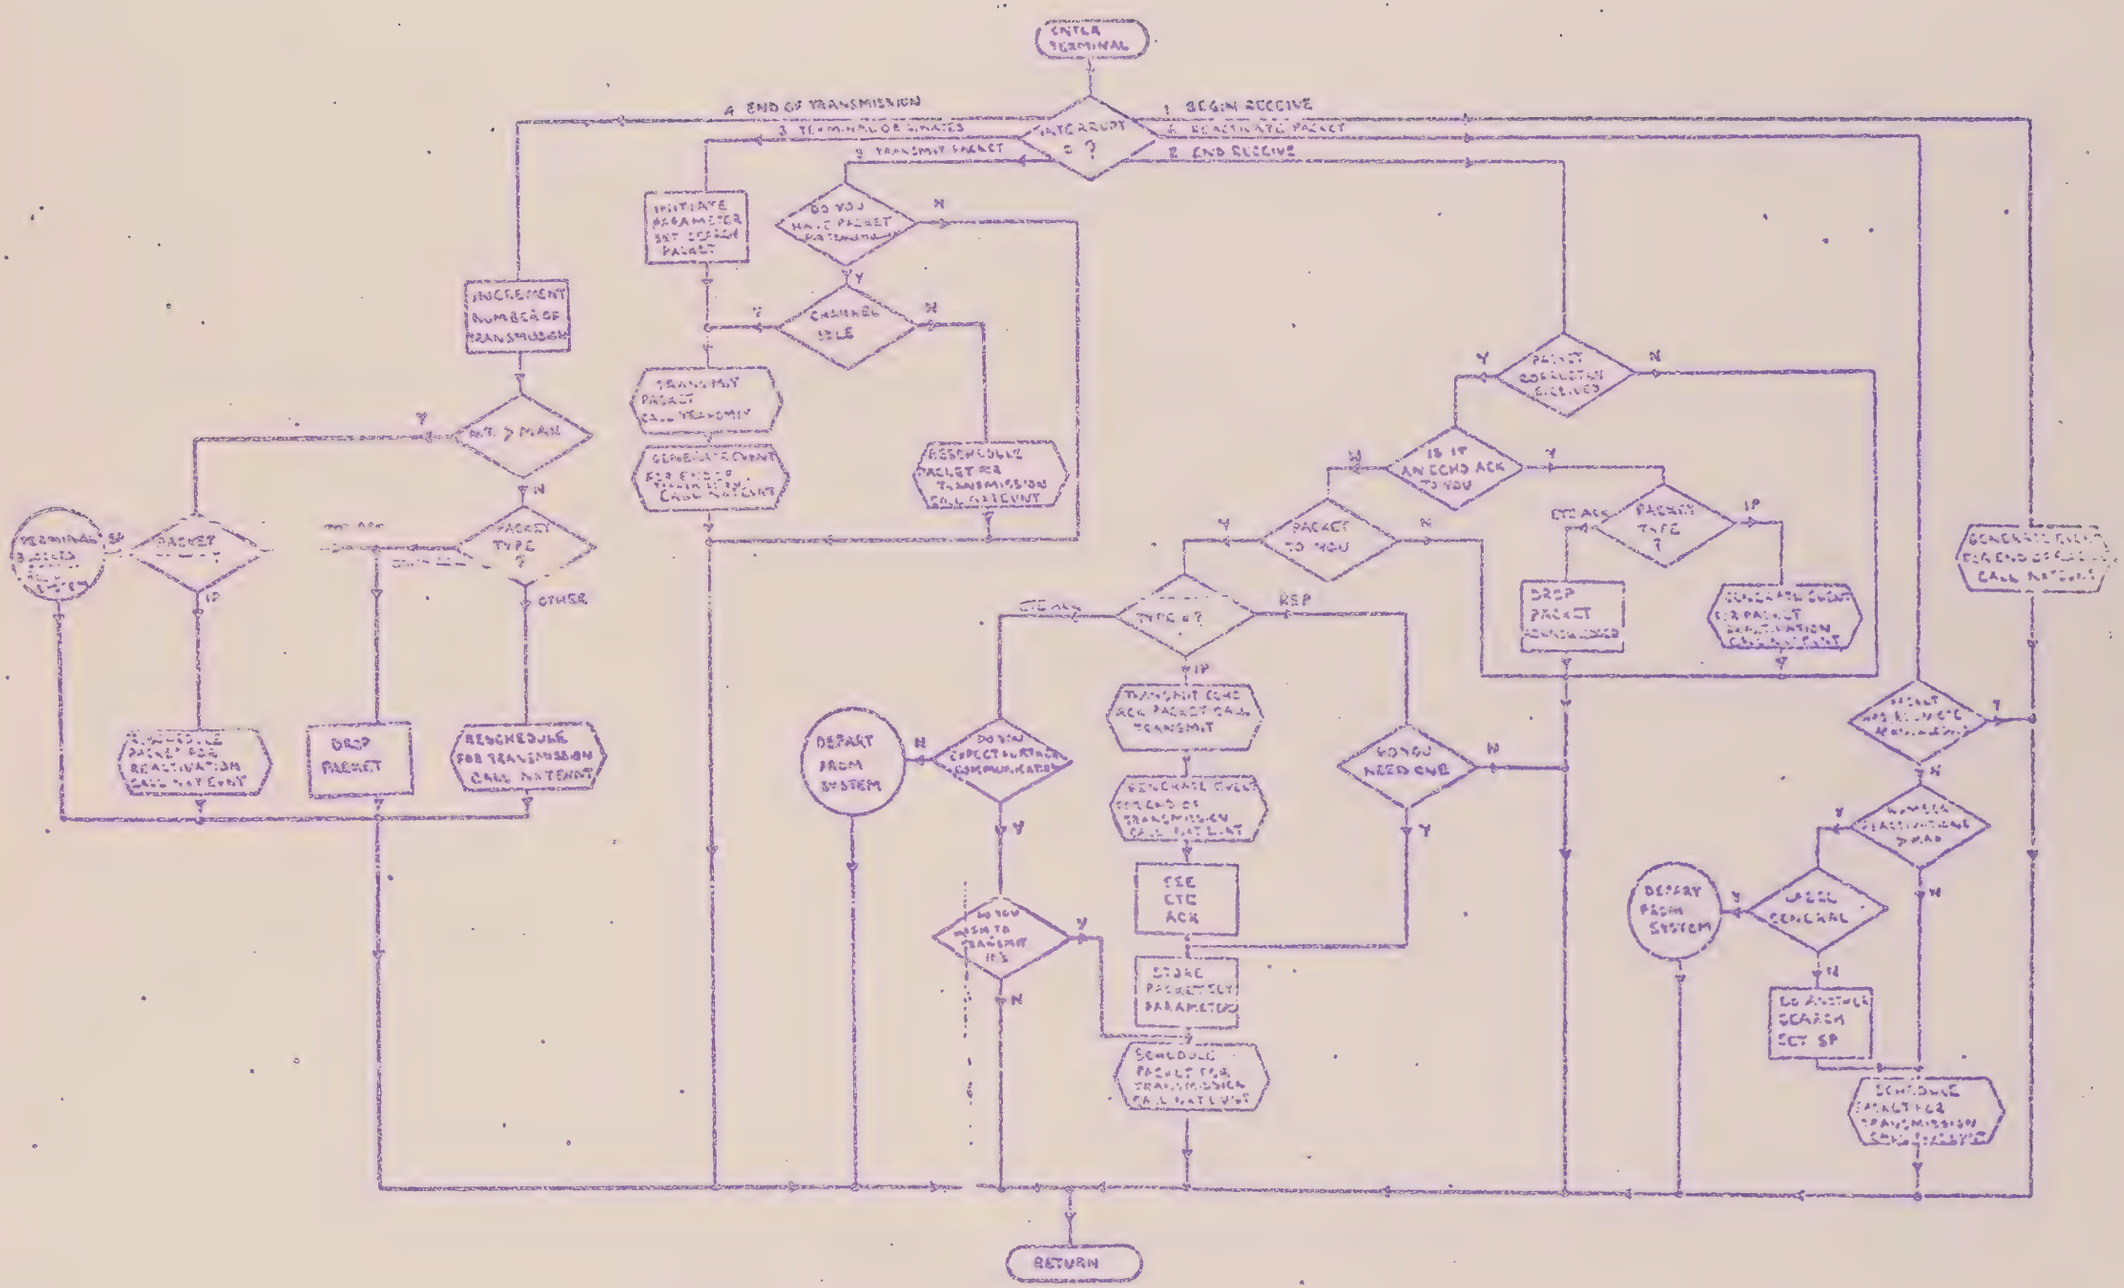
ENTER TERMINAL
INTERRUPT = ?
4 END OF TRANSMISSION
1 BEGIN RECEIVE
2 END RECEIVE
INCREMENT NUMBER OF TRANSMISSION
N.T. > MAX
PACKET TYPE ?
OTHER
DROP PACKET
RESCHEDULE FOR TRANSMISSION CALL NOTEVNT
DO YOU HAVE PACKET
CHANNEL IDLE
RESCHEDULE PACKET FOR TRANSMISSION CALL NOTEVNT
PACKET TO YOU
IS IT AN ECHO ACK TO YOU
PACKET TYPE ?
DROP PACKET
TYPE = ?
REP
DEPART FROM SYSTEM
DO YOU EXPECT FURTHER COMMUNICATION
DO YOU WISH TO TRANSMIT IT'S
SEE ETC ACK
STORE PACKET SET PARAMETERS
DO YOU NEED ONE
DEPART FROM SYSTEM
LABEL GENERAL
DO ANOTHER SEARCH SET SP
RETURN

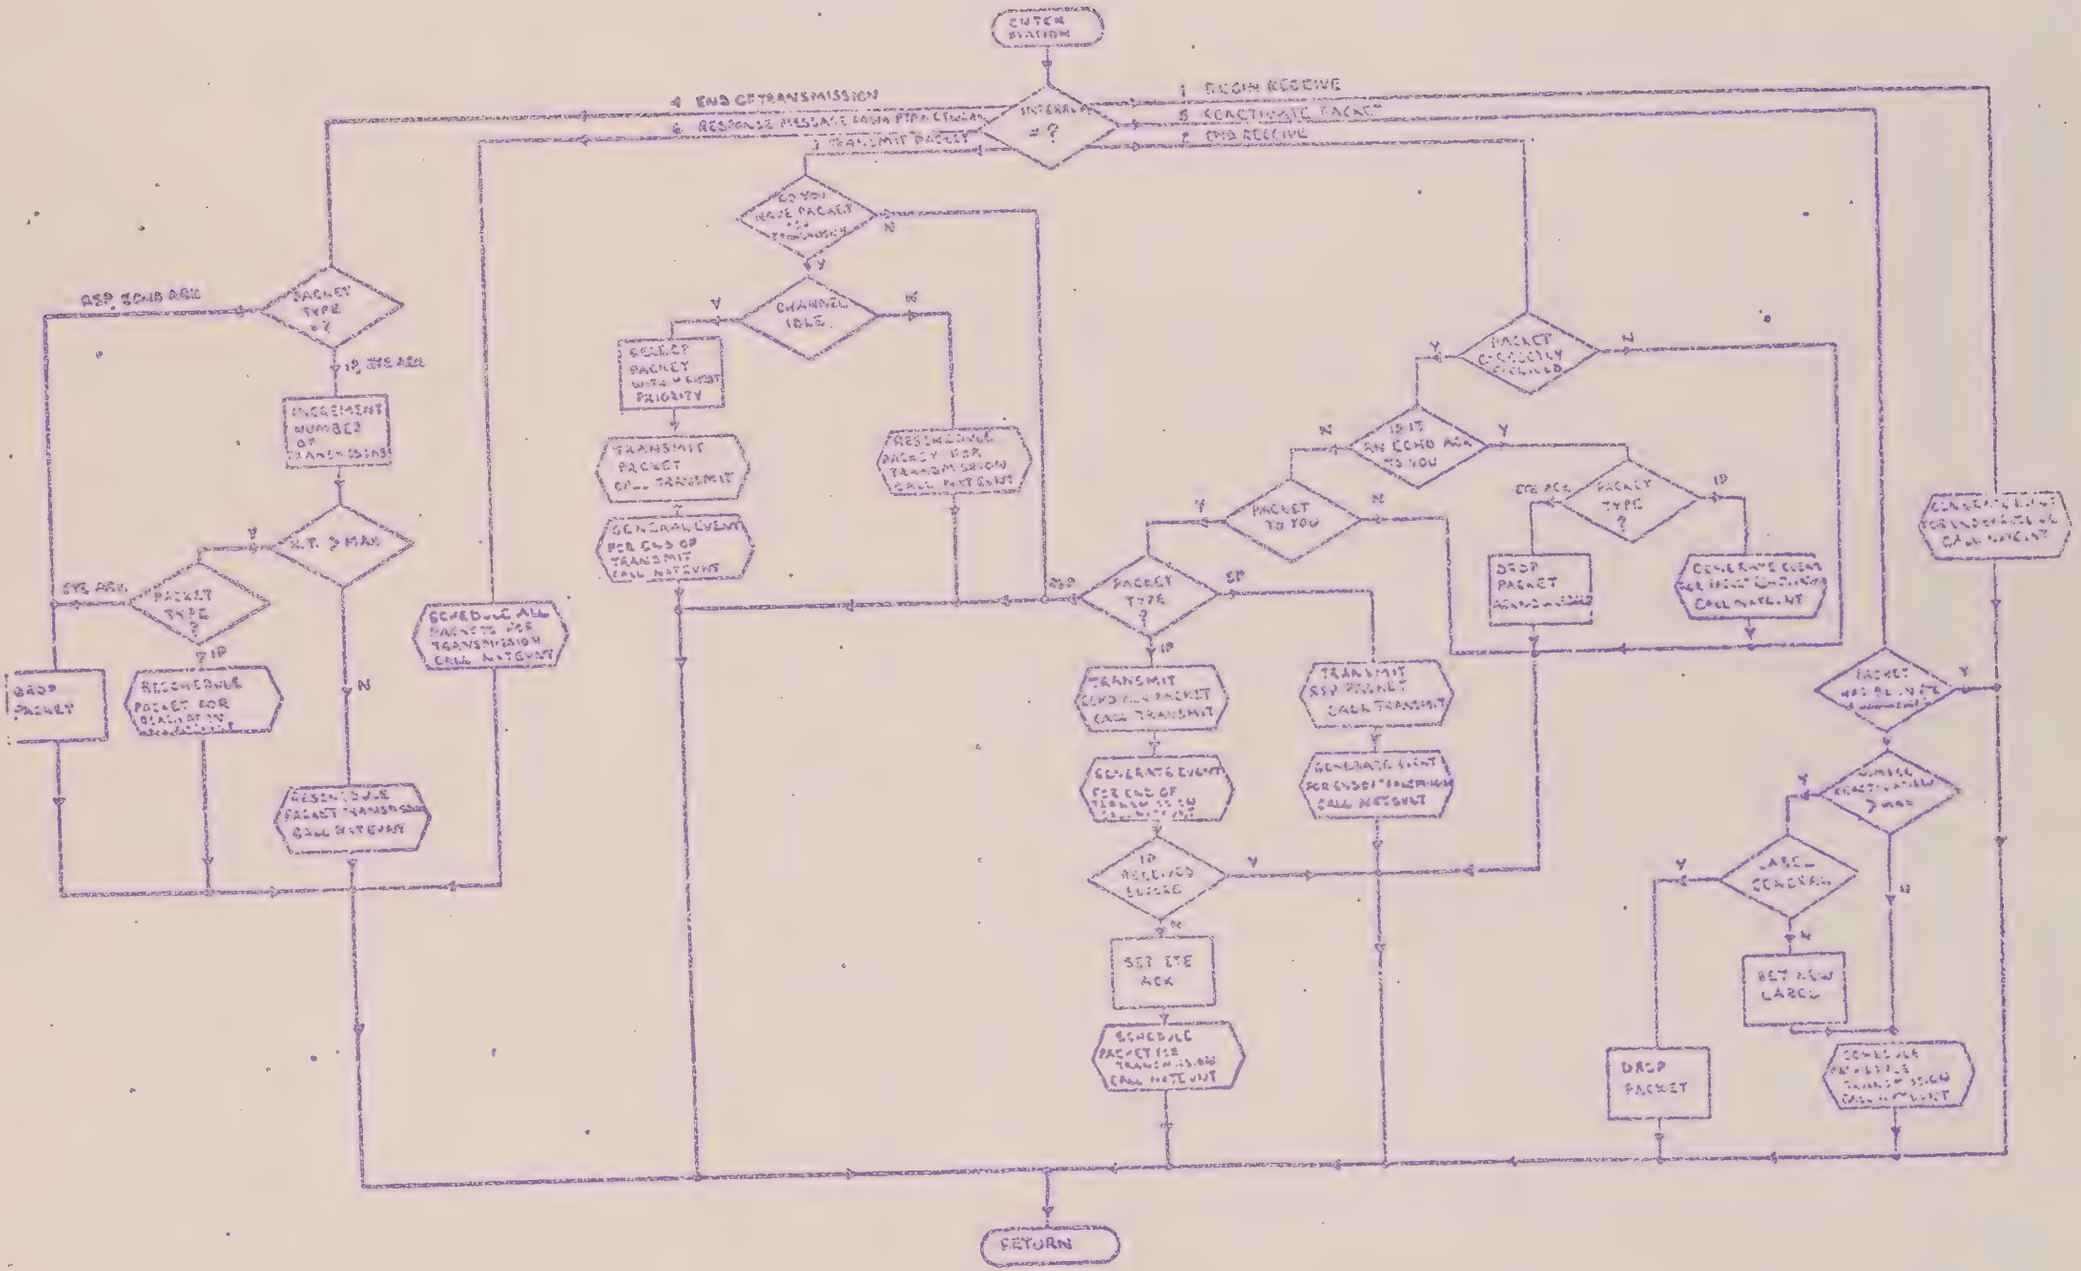

ENTER STATION
INTERRUPT = ?
1 BEGIN RECEIVE
2 END RECEIVE
3 TRANSMIT PACKET
4 END OF TRANSMISSION
5 REACTIVATE PACKET
6 RESPONSE MESSAGE FROM PTP ACTIVATION
PACKET TYPE = ?
RSP ECHO ACK
IP SYS ACK
INCREMENT NUMBER OF TRANSMISSIONS
N.T. > MAX
SYS ACK
PACKET TYPE ?
IP
DROP PACKET
RESCHEDULE PACKET FOR REACTIVATION
RESCHEDULE PACKET TRANSMISSION CALL NOTEVNT
SCHEDULE ALL PACKETS FOR TRANSMISSION CALL NOTEVNT
DO YOU HAVE PACKET FOR TRANSMISSION
CHANNEL IDLE
SELECT PACKET WITH HIGHEST PRIORITY
TRANSMIT PACKET CALL TRANSMIT
GENERAL EVENT FOR END OF TRANSMIT CALL NOTEVNT
RESCHEDULE PACKET FOR TRANSMISSION CALL NOTEVNT
PACKET CORRECTLY RECEIVED
IS IT AN ECHO ACK TO YOU
PACKET TO YOU
PACKET TYPE ?
SP
TRANSMIT ECHO ACK PACKET CALL TRANSMIT
GENERATE EVENT FOR END OF TRANSMISSION CALL NOTEVNT
IP RECEIVED BEFORE
SET ETE ACK
SCHEDULE PACKET FOR TRANSMISSION CALL NOTEVNT
TRANSMIT RSP PACKET CALL TRANSMIT
GENERATE EVENT FOR END OF TRANSMISSION CALL NOTEVNT
ETE ACK
PACKET TYPE ?
IP
DROP PACKET
GENERATE EVENT FOR END OF TRANSMISSION CALL NOTEVNT
PACKET HAS BEEN ETE
NUMBER REACTIVATIONS > MAX
LABEL GENERAL
BET NEW LABEL
DROP PACKET
GENERATE EVENT FOR END OF TRANSMISSION CALL NOTEVNT
RETURN

5. SUBROUTINES OF THE SIMULATOR

5.1 Data Structure and Management Subroutines

EVENT	Takes the next event out of the Event Data Structure for execution.
INHEAP	Adds an event to the heap in the Event Data Structure.
INMESS	Allows the introduction of special messages such as acknowledgments, control messages and the like into the Message Data Structure.
INPUT	Reads in the input parameters and determines the placing of repeaters and stations.
MESSREL	Is called by device routines to release a message as soon as all packets representing the message are deleted.
NEWMESS	Generates next exogenous message and adds to the Message Data Structure.
NEWPACK	Adds a new packet to the Packet Data Structure.
NXTEVNT	Adds a new event to the Event Data Structure.
OUT	Prints out intermediate data for debugging.
OUTHEAP	Takes the index of the next event time from the heap.
OUTPUT	Analyzes, and prints out the final statistics.
PRSIM	The driver routine. (main program)
MEASURE	Collects data on system performance.
MCOUNT	Counts the number of packets associated with each terminal which are stored in the system.

5.2 Communication and Device Subroutines

REPEAT	Main subroutine of repeater.
STATION	Main subroutine of station.
TERMINAL	Main subroutine of terminal.
DEVINIT	Reads parameters which define the particular communication system, labels, and flow control parameters. Initializes states of devices.

BGNRCV	Maintains states of devices related to the RF channel (e.g. number of overlapping packets)
ENDRCV	Same as above at the end of packet reception.
ECHO	Records that a device is receiving an echo acknowledgment.
SRER	Called when a non-overlapping packet is received for testing packet type and label.
ALTROUT	Called after repeater receives an RTS, checks whether repeater needs one, and has not used one before.
SRTTRT	Transmits packet and generates an event for the end of packet transmission.
REPNEXT	Determines which packet of a repeater is to be transmitted next.
TERSTOR	Stores correct IPs received by terminal and generates event for transmission of an ETE Ack.
SNREEKO	Called by station after receiving an echo Ack. Identifies and maintains the queue in which the acknowledged packet is stored. If it is an IP then it transfers the packet to another queue where it waits for an ETE Ack or for reactivation.
SHIFTQ	Shifts packets in the various queues of the station.
SNREPAK	Called by station after correctly receiving a packet. If the packet is an ETE Ack then subroutine drops the packets acknowledged, maintaines proper queues and the message counts. If it is an IP, subroutine verifies that same packet has not been received before and if so it generates packet and an event for transmission of an ETE Ack, and also generates a random time and an event for the arrival of the response message from the PTP network.

RESPONS	Called by station, sets all response packets to a terminal into the active queue and generates events for transmitting them.
SEKOTRT	Used by station to transmit an Echo Ack for the last hop.
CONNECT	Determins the most efficient repeater to which station should address packet when transmitting to a terminal.
TRANSMT	Called by a device which transmits a packet. Puts packet in list structure; determines all the devices that should receive the packet, the exact time for beggining to receive it; and generates the events to devices.

Summary of Acronyms

Ack	-	Acknowledgement
AR	-	Active Receive
ART	-	Active Receive and Transmit
AT	-	Active Transmit
ETE	-	End-to-end
HBH	-	Hop-by-hop
IP	-	Information packet
Label	-	An address assigned to a device for routing purposes
MHN	-	Maximum handover number
MNT	-	Maximum number of transmissions
PR	-	Passive Receive
PTP	-	Point-to-point
RSP	-	Response to search packet
SP	-	Search packet

References

1. A. Kershenbaum and R. VanSlyke, "Computing Minimum Spanning Trees Efficiently," Proceedings of the 1972 ACM Conference, Boston, August, 1972, pp. 518-527.

2. D. Knuth, *Sorting and Searching: The Art of Computer Programming, Vol. 3*, Addison-Wesley, 1973, Section 5.2.3.

3. J.W.J. Williams, "Heapsort," Algorithm 232, ACM Collected Algorithms, June, 1964.

4. NAC (a): "Packet Radio System-Network Considerations" ARPANET Network Information Center, Stanford Research Institute, Melno Park, California.

5. NAC (b): "Routing and Acknowledgement Schemes for the Packet Radio System", ARPANET Network Information Center, Stanford Research Institute, Melno Park, California.

6. NAC (c): "Packet Radio Simulation Results - I", ARPANET Network Information Center, Stanford Research Institute, Melno Park, California.

7. L. Kleinrock and F. Tobagi, "Carrier Sence Multiple Access (CSMA)", ARPANET Network Information Center, Stanford Research Institute, Melno Park, California.